Representation of natural reflectance spectra by auto-associative neural network

Huw C Owens, Laura Doherty & Stephen Westland
Colour & Imaging Institute
University of Derby
Derby

Abstract

The representation of spectral reflectance data by linear systems has been extensively studied [1-3] and it has been shown that colour information can be accurately represented by relatively few, although certainly more than three, parameters. In contrast Usui, Nakauchi and Nakano [4] have suggested that a non-linear system such as an auto-associative neural network can allow surface spectral reflectance data to be encoded by, and subsequently recovered from, just three parameters. We have repeated the analysis by Usui *et al.* using a set of 1269 Munsell reflectance spectra and have considered the representation of colour signals resulting from the spectral energy distributions of the Munsell surfaces viewed under D65 illumination. A five-layer wine-glass-shaped auto-associative neural network was used to encode and subsequently decode both reflectance spectra and colour signals. The middle layer of the neural network contained between 2 and 6 units so that the network was constrained to find efficient representations. Colour difference errors of reconstruction reduced with increasing number of hidden units as expected on the basis of theoretical considerations. Errors were smaller for networks trained with reflectance data rather than with colour signals derived from D65 illumination. Colour signals for surfaces viewed under D65 illumination are less constrained than the spectral reflectances of surfaces themselves. Spectral properties of light sources must be taken into account in computations of sampling rates required for recovery of colour signals and the subsequent recovery of surface reflectance spectra.

1. Introduction

In order to gain a deeper understanding of colour representation within the human visual system there have been several analyses of the reflectance spectra of Munsell colour chips. Munsell chips were selected and arranged on the basis of visual appearance criteria and thus their surface spectral reflectances may in some way reflect the colour representation of the human visual system. It is not at all clear however, that conclusions drawn from analyses carried out using the Munsell reflectance data can be extended to other types of surface [5]. Cohen (1964) first analysed the surface spectral reflectances of 433 randomly selected Munsell colour chips and conducted a principal component analysis by the Karhunen-Loeve (KL) expansion [6]. Cohen concluded that surface reflectance spectra can be described by a linear model using three or four parameters and discussed the results in terms of trichromatic colour vision. Maloney extended Cohen's analysis to 462 Munsell colour chips and concluded that the linear KL model fit the data when five to seven parameters are used [2]. Other workers have suggested that seven or eight linear parameters are required based on an analysis of 1257 Munsell chips [7-8].

One of the aims inherent in these studies was to find the most efficient basis set to represent sets of spectral reflectance data. If the results are compared with the physiological data there are similarities [e.g. 9] with the spectral properties of opponent channels. It has been suggested that the opponent coding of colours in the human visual system represents an efficient coding strategy [10]. The accuracy of the linear models is primarily dependent on the number of parameters allowed for the representation. Such analyses can therefore provide

important information about the potential for recovery of original signals from sensor devices (either biological or man-made) that capture signals using a small number of channels.

In contrast, some workers have considered the efficiency of representation of non-linear systems. For example, Usui *et al.* (1992) used an auto-associative five-layer neural network trained with a standard back-propagation learning algorithm to generate an identity mapping of surface-spectral reflectance data for 1280 Munsell colour chips [4]. The network could be considered to consist of encoding and decoding sub-networks. Once the identity map was established the response pattern of the middle layer was determined in relation to various Munsell colour chip inputs. It was concluded that three hidden units in the middle layer provided an optimum colour representation in the sense that higher order representations were less efficient. In addition, in the case of the three-dimensional representation, each of the three hidden units had some correspondence with a physiological colour channel; that is, luminance, red-green or blue-yellow.

Our study is concerned with the representation of natural and man-made reflectance spectra by linear and non-linear systems. Here we present some replication of the work of Usui's team and an extension of that work to the representation by autoassociative networks of colour signals from Munsell colour chips viewed under natural illumination.

2. Methods

A set of reflectance data were obtained which had been measured at 10nm intervals between 400nm and 700nm for 1269 Munsell surfaces [8,11]. A five-layer wine-glass-shaped auto-associative neural network (Fig. 1) was used to encode and subsequently decode the reflectance spectra of a randomly selected training set of 772 samples.

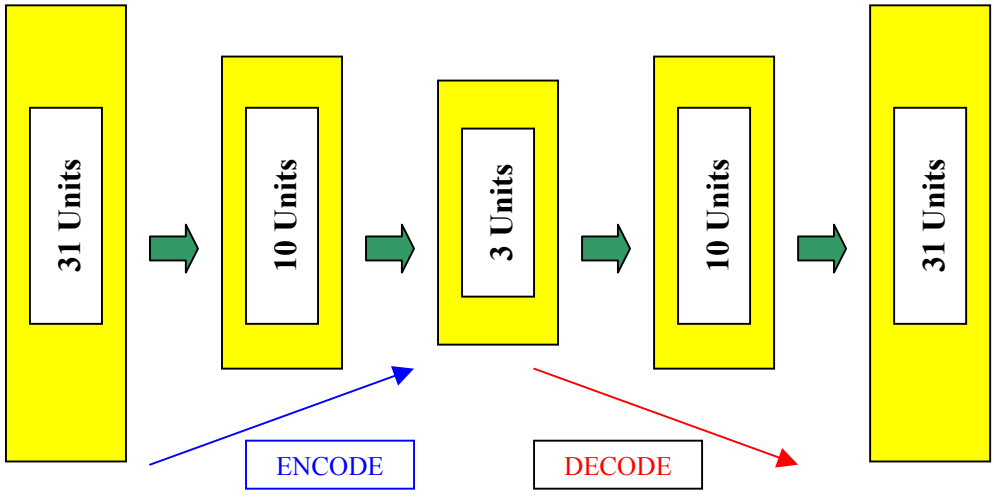

The yellow rectangles represent layers of processing units. The green arrows represent full forward connectivity between the layers. The direction of processing is from left to right. The leftmost (input) contains 31 units; each unit receives information at a wavelength between 400 and 700nm. The rightmost (output) layer similarly contains 31 units.

Figure 1: *Schematic diagram showing the neural network structure.*

The units in each layer received weighted input from each unit in the previous layer. Each unit (apart from those in the input layer) also received weighted input from a bias unit with fixed activation of 1; this is a standard feature of multi-layer perceptron networks. During training of the network the reflectance spectrum of each sample in the set was presented in turn to the input layer of the network. The activations of the units in the output layer were computed and compared with the reflectance spectrum that was presented at the input. Standard back-propagation algorithms [12] were used (implemented in MatLab) to find a set of weights that minimized the output errors for all samples in the training set. The network was forced to find an efficient representation of the colour data since the middle layer contained only 2-6 units but the reflectance spectrum was regenerated at the output layer. The selection of 10 hidden units in the 2^{nd} and 4^{th} layers of the network (Fig. 1) was arbitrarily chosen to be greater than 6 and less than 31. Networks were considered trained when the root mean squared (RMS) error between outputs and reflectance spectra, on a separately selected test set of 243 samples, was less than the RMS error for a linear coding of the data using six parameters.

Subsequently, networks were also trained to represent colour signals for the training set under D65. Colour signals were produced by multiplying normalised illuminants with the surface reflectance spectra. The normalised illuminant was given by dividing the original illuminant by the mean energy of the illuminant. We present results only for illuminant D65.

The performance of the trained networks was tested on the training set and on a separate validation set of 253 samples. The importance of the validation set is that it is not used, even for testing, at any stage during the development of the neural network training and therefore represents a statistically independent estimate of trained network performance. Performance was assessed in all cases by the mean CIELAB colour difference ΔE between actual and reconstructed spectra.

Reflectance spectra were also analysed by computing the principal components or basis functions of linear representations.

3. Results

Figure 2 shows reconstruction spectra for varying numbers (3-6) of hidden units for a typical reflectance spectrum. Figure 3 shows a similar reconstruction using varying numbers (1-10) of linear basis functions.

Networks with only 2 units in the middle layer could not be trained to achieve the desired RMS error. Table 1 therefore shows data obtained with 3, 4, 5, and 6 units in the middle layer. As expected, the CIELAB colour difference between the actual data and the data reconstructed from the efficient representations reduced as the number of hidden units increased. Colour differences were slightly higher when colour signals, rather than reflectance spectra, were encoded. Theoretical considerations support this finding since it has been postulated that the bandlimit of colour signals should be equal to the sum of the bandlimits of the respective surfaces and illuminant [2].

Table 2 shows data obtained by the principal-component analysis. As expected, the CIELAB colour differences between the actual spectra and the reconstructed spectra from the efficient representations reduced as the number of components used in the reconstruction increased.

	Mean CIELAB ΔE			
Hidden Units	Train set	Test set	Validation set	All data
3	8.28	8.92	8.44	8.44
4	4.28	4.49	4.44	4.35
5	5.27	6.04	5.76	5.51
6	3.7	3.93	3.71	3.75

(a) reflectance data

	Mean CIELAB ΔE			
Hidden Units	Train set	Test set	Validation set	All data
3	9.53	9.86	9.37	9.56
4	10.78	11.18	10.91	10.89
5	4.35	5.2	4.57	4.56
6	5.13	5.5	5.04	5.18

(b) colour signals

Table 1: *Mean CIELAB colour differences for networks trained with (a) reflectance data and (b) colour signals.*

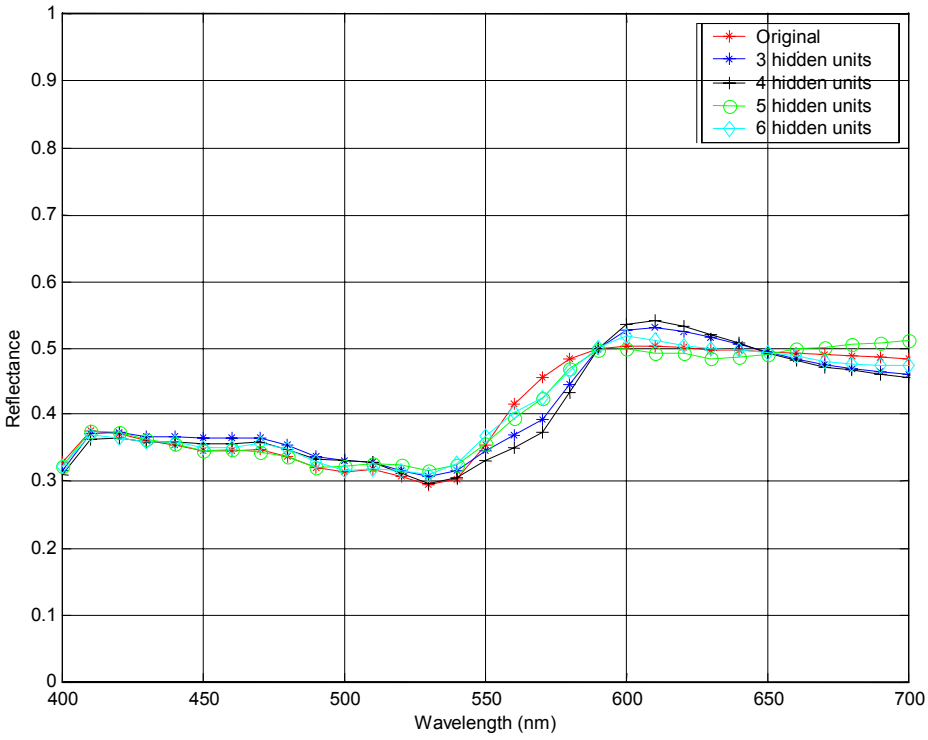

Figure 2: *Reconstruction of a typical Munsell spectrum using an auto-associative neural network*

	Mean CIELAB ΔE			
PCA component	**Test set**	**Validation set**	**Training set**	**Test set data**
1	23.23	23.79	22.79	23.08
2	17.47	16.73	16.15	16.52
3	3.50	3.02	2.90	3.04
4	2.56	2.40	2.32	2.38
5	1.30	1.30	1.30	1.30
6	0.31	0.29	0.28	0.29
7	0.27	0.25	0.25	0.25
8	0.20	0.19	0.18	0.19
9	0.16	0.17	0.15	0.16
10	0.11	0.11	0.10	0.10

Table 2: *Mean CIELAB colour differences using n principal components in reconstruction of reflectance spectra.*

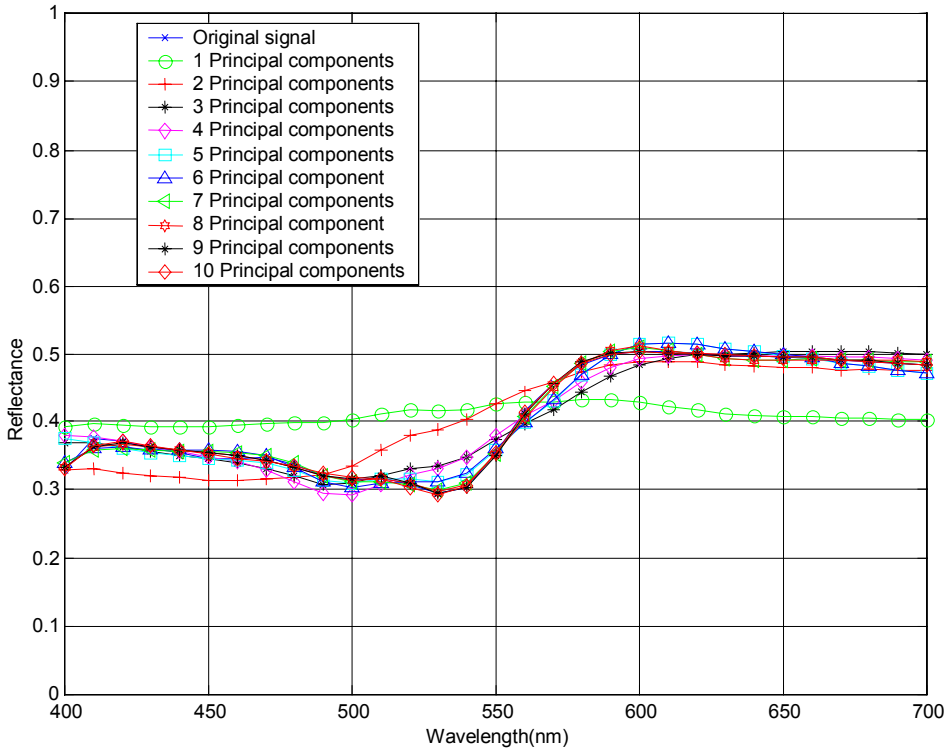

Figure 2: *Reconstruction of a typical Munsell spectrum using a linear system.*

Interestingly, if a linear system is used six components are required to produce colour differences acceptable for many industrial applications (ΔE<1.0). Paradoxically, six parameters in the non-linear neural network allowed recovery to an accuracy on average of about 3.5 CIELAB ΔE units. It is, perhaps, surprising that the linear system outperforms the

non-linear system given the same number of degrees of freedom. One possible explanation for this is that the MSE error term that is minimised by both techniques is not appropriate if the goal is to minimise perceptual colour differences. Unlike colour differences, the MSE term does not take into account the spectral sensitivity of the human visual system.

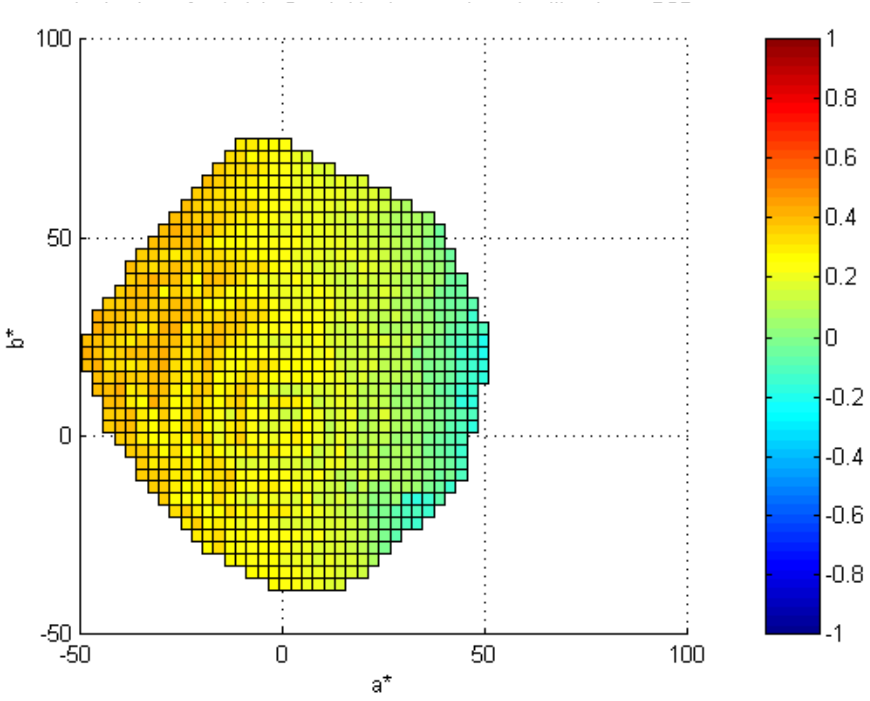

Figure 3: *Activation map for one of the middle units in a neural network trained possessing three middle units and trained with colour signals.*

There is some evidence that the neural network that was trained with three units in the hidden layer was able to learn that an opponent representation of the reflectance spectra and colour signals is most efficient. For example, Fig. 3 shows an activation map for one of the units in the middle layer of a network that had three units in the middle layer and that was trained using colour signals of Munsell chips under D65. The activation map was obtained by presenting colour signals to the input of the trained network and noting the activation of the middle units. Fig. 3 shows how these activations varied with the CIELAB *a* b** coordinates of the colour signals. The particular unit illustrated demonstrates tuning in the *a** direction. Other units were found that were mainly sensitive to luminance or to *b**.

4. Conclusions

We have trained auto-associative neural networks to find efficient representations of colour information for reflectance data and colour signals. Representations were more efficient for raw reflectance data than for colour signals under D65. Colour signals for surfaces viewed under D65 illumination are less constrained than the spectral reflectances of surfaces

themselves. The spectral properties of light sources must be taken into account in computations of sampling rates required for recovery of colour signals and subsequently reflectance spectra.

A neural representation using three units in the middle layer exhibits recovery errors equivalent to an average colour difference ΔE of over 9.0 CIELAB units. The representations become more efficient as the number of units in the middle layer is increased but even the network with six units in the middle layer produces average colour differences of 3.7 CIELAB units.

The fact that efficient representations cannot be found with only two hidden units is consistent with the findings of Usui *et al.* [4]. Furthermore the activation patterns for the network with three hidden units, in particular, share some properties with the opponent channels in the human visual system and with the principal-component or subspace representations of reflectance spectra reported by several workers [9,11]. Thus, opponent representations are highly efficient representations of colour information [10]. We find no evidence, however, that reflectance spectra, nor indeed colour signals, are inherently three-dimensional despite the fact that three-dimensional systems can represent a surprisingly large proportion of the variance in the data [4].

We were surprised that the reconstruction errors for the non-linear systems were greater than for the linear systems given the same number of degrees of freedom. This result could indicate that the neural network systems were not trained sufficiently. However, the neural network representations were trained to achieve a mean squared error equivalent to that of six linear components. Therefore the larger CIELAB colour differences shown for the neural network representation may simply indicate that MSE is not an appropriate measure to minimise since it does not take into account the spectral sensitivity of the human visual system.

Further work is underway to investigate the effect of other natural illuminants and of artificial light sources on the statistics of surface colour signals. These statistics will be characterised using linear and non-linear models and have important implications for the validity of computational models of colour vision (e.g. colour constancy [13] and transparency [14]) and sampling errors for machine-vision systems.

References

1. Maloney L.M. and Wandell B.A. (1986). *J. Opt. Soc. Am.*, A **3**, 3, 29-33.
2. Maloney L.M. (1986), "Evaluation of linear models of surface spectral reflectance with small numbers of parameters" *J. Opt. Soc. Am.*, A **3**, 1673-1683.
3. Marrimont D.A. and Wandell B.A. (1992), "Linear models of surface and illuminant spectra" *J. Opt. Soc. Am.*, A **9**, 11, 1905-1913.
4. Usui S., Nakauchi S. and Nakano M. (1992). "Reconstruction of Munsell color space by a five-layer neural network" *J. Opt. Soc. Am.*, A **9**, 4, 516-520.
5. Barlow H.B. (1982). "What causes trichromacy? A theoretical analysis using comb-filtered spectra" *Vision Res.*, **22**, 635-643.
6. Cohen J. (1964). "Dependency of the spectral reflectance curves of the Munsell color chips" *Psychon. Sci.*, **1**, 369-370.
7. Sobagaki (1984), Bull Electrotechnical Lab. Jpn. 48, 785-792.
8. Parkkinen J.P.S., Hallikainen J. and Jaaskelainen T. (1989), "Characteristic spectra of Munsell colors", *J. Opt. Soc. Am.*, A **6**, 318-322.
9. Skelton H. (1997). *PhD Thesis*, Leeds University.

10. Buchsbaum G. and Gottschalk A. (1983). "Trichromacy, opponent colours coding and optimum colour information transmission in the retina", *J. Opt. Soc. Am.*, B **220**, 89-113.
11. Jaaskelainen T., Parkkinen J.P.S. and Toyooka S. (1990), "Vector subspace model for color representation", *J. Opt. Soc. Am.*, A **7**, 725-730.
12. Fausett L. (1994), "Fundamentals of neural networks: architectures, algorithms, and applications", Prentice Hall.
13. Foster D.H. and Nascimento S.M.C. (1994), "Relational colour constancy from invariant cone-excitation ratios", *Proc. R. Soc. Lon.*, B **257**, 115-121.
14. Westland S. and Ripamonti C. (2000), "Invariant cone-excitation ratios may predict transparency", *J. Opt. Soc. Am.*, A **15**, 255-264.